Simple Auctions*

Olivier Compte
PSE, Paris

Andrew Postlewaite
University of Pennsylvania

First version: April 2010
This version: May 2012

Abstract

Standard Bayesian models assume agents know and fully exploit prior distributions over types. We are interested in modeling agents that lack detailed knowledge of prior distributions, and we achieve that putting restrictions on the strategies that agents are allowed to compare.

In auctions, that agents know priors has two consequences: (i) signals about own valuation come with precise inference about signals received by others; (ii) noisier estimates translate into more weight put on priors.

We revisit classic questions in auction theory, exploring environments in which no such complex inferences are made or allowed. The contribution is to shed a new light of classic insights; to provide (and advocate) more parsimonious and cognitively less demanding models of auctions; and to question some standard intuitions, in particular those based on the effect of information, i.e. the fact that in Bayesian models, no or poor information generally translates into concentrated posteriors.

*Compte: Paris School of Economics, 48 Bd Jourdan, 75014 Paris (e-mail: compte@enpc.fr); Postlewaite: Department of Economics, 3718 Locust Walk, University of Pennsylvania, Philadelphia, PA 19104-6207 (e-mail: apostlew@econ.upenn.edu). A large part of this work was done while Postlewaite visited the Paris School of Economics; their hospitality is gratefully acknowledged. We thank Paul Klemperer and Steve Matthews for helpful discussions. The authors thank the Gould Foundation for financial support, and Postlewaite thanks the National Science Foundation for financial support.

1 Introduction

I might go to an auction at Sotheby's, see a painting and decide that my value for a painting is \$1100, but have little idea of where my valuation stands in relation to the valuations of other bidders. That is, I might think it is equally likely that my valuation is the highest, second highest, or lowest; in other words, my valuation gives me little guidance in predicting my rank in the valuations. Had my value of the painting been \$1200 rather than \$1100 I might not think it any more likely that my valuation was the highest: whatever made my value higher might well have made other bidders' values higher as well, and I might find it impossible to sort out whether the basis of my higher value was common to all bidders or was idiosyncratic to me.

We are interested in considering auction models that reflect the above difficulty – that of using personal valuation as an instrument in predicting rank. Of course, some signals might be good instruments: If I see another participant at the painting auction arriving with a \$4000 suit, that may be a strong indication that my rank is low. Our view however is that personal valuation is often not a useful instrument to estimate rank, and our perspective is that a plausible model of auction should reflect that.

This is a change in perspective because standard auction theory implicitly poses no limits on the agent's ability to use personal valuation as an instrument to assess how one's value compares to others, or to disentangle common and idiosyncratic elements, resulting in behavior that is finely tuned the particular joint distribution over valuations that the analyst assumes.

Our approach is similar to that explored in Compte and Postlewaite (2012). In that paper, the agent gets an estimate of some underlying state s, say $x = s+\varepsilon$. Ideally, the agent would like to know his estimation error ε, but we should not expect him to be able to make precise inferences about the size or sign of ε based on the estimate x he forms. There too, we took the perspective that a plausible model of decision making should reflect that difficulty, contrasted with standard models that implicitly presume agents use very finely detailed knowledge of the joint distribution over state and estimate to make inferences about their current estimation errors.

A bidder participating in auctions faces a similar difficulty. He has a valuation v_i for the object and a possibly noisy estimate x_i of that valuation. In a first price auction, he wins if and only if his bid b_i exceeds some price p, and then earns $v_i - b_i$. There is an underlying process that defines the connection between the problem faced (v_i, p) and the data that the agent uses (x_i), characterized by a joint distribution ω over (v_i, p, x_i). The standard approach assumes the bidder knows the joint distribution ω, or behaves optimally as if he knew ω with great precision. We call this the *omniscient* perspective.

We propose that a bidder's behavior be driven by own welfare considerations, but only to a limited extent. We suggest intuitively plausible rules that an agent may employ given the type of data available to him, and we assume that among these plausible rules, the agent knows which is optimal. We thus depart from the standard/omniscient perspective in that we put restrictions on the set of rules that one compares. Without limitations, the agent's decisions would be the same as if he were omniscient.

This paper is part of a broader project that attempts to model agents that *lack detailed knowledge of prior distributions*. Following Wilson (1987)'s critique, a number of authors have questioned the omniscient perspective that agents would have precise knowledge of prior distributions. Attempts to weaken that perspective generally follow the traditional Bayesian route that takes anything that is not known as a random variable, assuming a known prior over priors and enriching the type's space with additional private signals.[1] Strategy restrictions constitute another route in that they have the effect of preventing agents from finely adjusting behavior to environment details that they cannot plausibly know.

An alternative interpretation of the model is that we examine the bidding behavior of *boundedly rational agents* and explore, within that framework, some of the standard issues addressed by the auction literature: existence, revenue rankings, and comparative statics. Compared to more traditional bounded rationality models, that generally assume that agents are making evaluation errors in comparing alternatives, we assume that agents do not make errors in comparing bid functions.[2] They are just not allowed to compare all bid functions.

The main contribution of the paper is to point out that the standard assumption that agents know and fully exploit prior distributions is not innocuous, either in terms of predictions or in terms of how our intuitions are shaped. That agents know priors has two consequences: (i) signals received come with a precise inference about signals received by others; (ii) noisier estimates translate into more weight put on priors;

(i) implies that shading behavior depends finely on one's inference about others' types – with higher types trying to extract higher rents. Revenue comparisons between auction formats depend finely on theses inferences, and comparative statics depend finely on how further information provided affect these inferences. A contribution of this paper is to provide a new perspective on these comparisons by exploring environments in which such complex inferences are not made or allowed.

(ii) implies that a more "poorly informed" bidder (i.e. with noisy value estimates) has more concentrated posterior beliefs. From the perspective of

[1] This is the route taken by the robustness in mechanism design literature (Bergemann Morris, 2005), as well as the global game literature (See Morris and Shin, 2006 for a survey). The underlying motivation is to avoid predictions that would be too sensitive to common knowledge assumptions (as in Rubinstein 1989's email game). The technique consists of "relaxing" common knowledge by adding further private signals (In the context of auctions, see Fang and Morris 2006, where the additional signals permit to break the one-to one relationship between value and beliefs).

These papers however assume common knowledge of the distribution over this enriched type space, making agent's optimization problem ever more sophisticated.

[2] The usual bounded rationality route consists in assuming one particular choice rule – in general subjective expected utility maximization, based on possibly mistaken evaluations. The source of the errors may be left exogenous (Block and Marchak, 1960), procedural (Osborne Rubinstein, 1998), or stem from the way beliefs are formed (Geanakoplos, 1989, Gilboa Schmeidler, 1995, Jehiel, 2005).

other players, his behavior is thus more predictable (unless he uses a mixed strategy). The consequences are numerous: rents are easier to extract from poorly informed agents; Ignorance promotes competition (between few symmetric agents) (Ganuza, 2004); In competing with a more informed agent, the less informed agent gets no rents (Milgrom, 1981); it also implies that for an uninformed agent, one benefit of (publicly) getting more precise information is to make behavior less predictable. For sellers, incentives to provide information thus have to balance efficiency gains with the additional rents that less predictable behavior generates (Bergemann Pesendorfer, 2007).[3]

One contribution of this paper is to question these insights, and point out that they rely on our standard modelling choice that assumes known priors. In many problems, there is no obvious reference point that would justify that a more poorly informed agent has a more predictable behavior. In the absence of such a reference point, or if agents lack detailed knowledge of the joint distribution over value and value estimates, it is likely that the opposite will be true, that is, that from the perspective of other players, behavior of more poorly informed agents is *less* predictable. Our modeling choices will reflect that possibility, allowing us to focus on what we think is the primary benefit of getting better signals, that of reducing the discrepancy between value and value estimates.

Finally, another contribution of the paper is to promote theory that takes an agent's perspective rather than an analyst's perspective, starting from the *cognitively* less demanding models, and then possibly asking what further sophistication brings about. For example, to the analyst or to an omniscient bidder, the independent private value model seems relatively simple. Cognitively however, the model is quite demanding: it requires the agent to combine in subtle ways information about rank and dispersion based on each possible value realization.

Section 2 starts by describing the bidder's *decision problem*, and introduces our main modelling assumption (strategy restrictions) without reference to the specific auction problem that the agent will ultimately face. Section 3 moves to the analysis of a simple first price auction with private values. Existence, revenue comparisons, and comparative statics results are obtained. Section 4 incorporate additional private signals into bidding behavior, showing that providing rank information has an ambiguous effect on sellers' revenue. Section 5 introduces noisy estimates. Section 6 provides a discussion and relates results to the literature. Section 7 concludes. Appendix A and B extends the analysis to several other standard questions: buyer-seller relationships, and comparison of discriminatory and uniform auctions (with unitary demands).[4]

[3] These remarks apply more generally, to other strategic situations. In Kamenica and Gentzkow (2011) for example, one benefit from information transmission is the dispersion in posteriors that information transmission generates, explaining why sender does not benefit from persuasion when his value function (expressed as a function of the receiver's beliefs) is concave.

[4] The working paper (Compte and Postlewaite 2010) includes extensions to asymmetric auctions, auctions for bundles and sequential auctions.

2 The bidder's decision problem

We consider first price auctions and start by describing the *decision problem* that a given agent faces. In any first price auction, the agent chooses a bid, denoted b. That agent wins if and only if b exceeds some price p. Letting v denote his valuation for the object being sold, he obtains $v - p$ if he wins, and 0 otherwise. So the agent's preferences over his possible alternatives (the bids b) depend on the *state* (v, p), and they are characterized by:

$$\begin{aligned} u(b, v, p) &= v - p \text{ if } b \geq p \\ &= 0 \text{ otherwise.} \end{aligned}$$

The agent however faces some uncertainty: he does not know the state (v, p) but has some imperfect knowledge of it. This is modelled by assuming that he gets a signal or data $z \in Z$ correlated with the true state (v, p). The joint distribution over state and signal is denoted ω.[5]

The "data" z may take various forms. Our basic auction model will assume private values, that is, $z = v$. More generally the data z may consist of a noisy estimate of v. Formally, we define

$$x = v + \varepsilon$$

where ε is a noise term independent of v, and we set $z = x$.[6]

The standard approach

Rational decision making consists of assuming that for any signal z that he might receive, the agent chooses the alternative that maximizes his expected welfare. Denoting by $r^*(z)$ that optimal decision, we have:

$$r^*(z) \equiv \arg\max_b E_\omega[u(b, v, p) \mid z].$$

Alternatively, one may define the set $\overline{R}$ consisting of all possible decision rules $r(\cdot)$ that map Z into A. For each rule $r \in \overline{R}$, one can define the expected utility (or performance)

$$v_\omega(r) = E_\omega u(r(z), s).$$

The optimal decision rule r^* solves:

$$r^* = \arg\max_{r \in \overline{R}} v_\omega(r).$$

An alternative approach.

[5] When analyzing an auction with several strategic bidders, the distribution ω is endogenous. For now however, we shall keep it exogenous.

[6] In Section 4, we shall consider the case where bidders receive some signal correlated with rank (i.e. the $4000 suit), that is

$$z = (v, \rho)$$

where $\rho \in \{H, L\}$ is a signal (High or Low) correlated with p.

Our perspective is that it is not plausible that the bidder would know ω, or would behave as if he knew ω. We propose an alternative approach, based on the idea that the agent's behavior is driven by welfare consideration, but only to a limited extent. Rather than following r^*, the agent follows a rule that is optimal within a more limited set of rules, denoted R. Formally, we let $r^*_\omega = \arg\max_{r \in R} v_\omega(r)$. We emphasize that we do *not* conceive of the agent carrying out this maximization; rather the agent restricts attention to a set of rules and has learned or been taught the circumstances under which each of the rules is optimal.

> **Main assumption (A1)**: The set of considered rules R consists of a limited set, i.e., $R \subset \overline{R}$. Although the agent does not know ω, he identifies (or learns) which rule r^*_ω is optimal in R, and he follows it.

The standard approach corresponds to the case where no restrictions are imposed on R, i.e., $R = \overline{R}$. *A*1 thus corresponds to a weakening of the standard assumption. In what follows, we shall propose a class of plausible rules.

Plausible bidding rules. The data that the agent gets is z, and we consider a family of bidding rules that depend on z. We consider the case where z is a noisy estimate of v, and for any γ, we define rule r_γ as:

$$r_\gamma(z) \equiv z - \gamma.$$

The parameter γ characterizes how *unaggressive* the agent is (with higher γ implementing *lower* aggressiveness). It may also be interpreted as a level of cautiousness (if the agent fears that $z > v$). Obviously, one could imagine other ways to parameterize aggressiveness or cautiousness, and we certainly do not wish to argue in favor of a specific shape. Our main motivation lies not in the shape of the rules considered, but in the assumption that the agent is unable to adjust γ to each particular realization of z for a class of problems, reflecting the idea that z is not used as an instrument to adjust how cautious or aggressive one ought to be. Throughout most of the paper, we shall assume that

$$R = \{r_\gamma\}_{\gamma \in \mathcal{R}}.$$

Optimal bidding rules.

The optimal bidding rule is characterized by a single parameter γ^* that measures the optimal extent of shading. We illustrate with two cases.

(i) Private values ($z = v$).

Define $\phi(\gamma)$ as the probability of winning when the agent uses r_γ:

$$\phi(\gamma) \equiv \Pr(v - p > \gamma).$$

The expected payoff that the agent gets from using rule r_γ is:

$$v(r_\gamma) = \gamma\phi(\gamma).$$

The derivation of the optimal bidding rule is thus analogous to a standard monopoly pricing model in which $\phi(\gamma)$ is interpreted as a demand function. Optimal shading γ^* is characterized by the following first order condition:[7]

$$\gamma^* = \frac{\phi(\gamma^*)}{-\phi'(\gamma^*)},$$

Shading is thus larger when the agent has higher chances of winning (that is, if $\phi(\gamma)$ shifts up), or when the distribution over $v - p$ is more dispersed (that is when $\mid \phi' \mid$ shifts down). The latter effect can be interpreted as a Bertrand competition effect: it is stronger when small changes in shading induce large changes in the probability of winning.

(ii) Noisy estimates ($z = v + \varepsilon$).

Now define

$$\phi_\varepsilon(\gamma) \equiv \Pr(v + \varepsilon - p > \gamma) \text{ and } \psi_\varepsilon(\gamma) \equiv E[\varepsilon \mid v - p + \varepsilon = \gamma].$$

$\phi_\varepsilon(\gamma)$ is the probability of winning when the agent uses r_γ, while $\psi_\varepsilon(\gamma)$ is the expectation of the estimation error conditional on using r_γ and winning by a 0-margin. The optimal bidding rule is characterized by a shading level γ^* that solves:

$$\gamma^* = \frac{\phi_\varepsilon(\gamma^*)}{-\phi'_\varepsilon(\gamma^*)} + \psi_\varepsilon(\gamma^*)$$

In words, optimal shading is now derived from two considerations. First, as in the private value case, it depends on demand elasticity. In addition, when estimates are noisy, a bidder runs the risk of getting the object only because the error term ε was positive (and high). In other words, *noisy estimates* make a bidder subject to *a winner's curse, or selection bias*, and an optimal reaction to that possibility is cautiousness – or shading more.[8]

Due to noisier estimates, optimal shading may thus increase for two different reasons: more dispersion in estimates (thus weakening the Bertrand competition effect) and a stronger winner's curse effect.

3 A basic auction model.

We now apply our approach to strategic interactions between n bidders. We assume that bidder i's value for an object is of the form

$$v_i = \alpha + \theta_i,$$

where α represents a common component of all bidders' values and θ_i an idiosyncratic component of i's value. This model captures the logic of the painting example: i may know his value v_i but does not know how much of his value

[7] Assuming that $\gamma + \frac{\phi(\gamma)}{\phi'(\gamma)}$ is increasing, the first order condition is sufficient.

[8] Note that the sign of the term $\psi_\varepsilon(\gamma)$ depends on the distribution over $v - p$. With fierce enough competition, prices tends to be high and $\psi_\varepsilon(\gamma)$ positive.

is idiosyncratic, that is, he does not observe either α or θ_i. We assume that the vector of idiosyncratic terms is drawn independently of α, from some joint distribution g.

Each bidder faces a decision problem similar to that described in Section 2. For each bidder i, we consider a set of bidding rules R_i, where each rule $r_i \in R_i$ maps the data z_i that the bidder receives into a bid b_i. Throughout the paper, we shall mostly focus on the case where z_i is a possibly noisy estimate of v_i, that is,

$$z_i = v_i + \varepsilon_i,$$

where error terms ε_i are drawn independently of v_i. In Section 4, we shall also consider the case where z_i includes information about rank. For now, we assume as before that

$$r_\gamma(z_i) = z_i - \gamma \text{ and } R_i = \{r_\gamma\}_{\gamma \in \mathcal{R}}$$

For any vector $\gamma = (\gamma_1, ..., \gamma_n)$, each rule profile $r_\gamma = (r_{\gamma_1}, ..., r_{\gamma_n})$ induces an expected payoff which we denote by $v_i(\gamma_1, ..., \gamma_n)$. An equilibrium is then defined in the usual way.

> **Definition:** *A shading vector* $\gamma^* = (\gamma_1^*, ..., \gamma_n^*)$ *is an equilibrium iff* $v_i(\gamma^*) \geq v_i(\gamma_i, \gamma_{-i}^*)\ \forall \gamma_i$.

Compared to standard auction models, behavior is characterized by a one-dimensional parameter, and equilibrium behavior is relatively simple to derive. It has the same complexity as a problem of price competition with differentiated products, where γ_i is the price set by player i and $v_i(\gamma_i, \gamma_{-i})$ is the profit that results from the price vector (γ_i, γ_{-i}).

We illustrate the approach with the private value case.

3.1 Private values

We assume that each bidder observes his valuation without noise, and that the vector of idiosyncratic terms is drawn from a symmetric distribution. We look for a symmetric equilibrium in which all bidders pick the same rule r_{γ^*}.

Define $\phi(y)$ as the probability that bidder i's valuation exceeds all others by at least y:

$$\phi(y) = \Pr(\theta_i - \max_{j \neq i} \theta_j > y).$$

We have

$$v_i(\gamma_i, \gamma^*) = \gamma_i \phi(\gamma_i - \gamma^*).$$

The first order condition for a symmetric equilibrium can thus be written as[9]

$$\gamma \phi'(0) + \phi(0) = 0.$$

By symmetry, $\phi(0) = \frac{1}{n}$, thus implying the following Proposition:

[9] This first order condition is sufficient when $y \to y + \frac{\phi(y)}{\phi'(y)}$ is increasing in y. We shall come back to existence issues in the next Section.

Proposition 1: *In a symmetric equilibrium, we must have:* $\gamma^* = \frac{1}{-\phi'(0)n}$.

In other words, equilibrium shading is driven by the expected chance of winning ($1/n$) and the dispersion of the idiosyncratic terms. Indeed, to interpret $-\phi'(0)$, consider the case where the idiosyncratic elements θ_i are i.i.d., drawn from a distribution with density f. Let $\theta^{(2)} = \max_{j \neq i} \theta_j$. It is easy to check that

$$-\phi'(0) = Ef(\theta^{(2)}).$$

The coefficient $-1/\phi'(0)$ thus measures, from the perspective of the winner, the dispersion of second highest valuations. In the special case where the distribution is uniform on the interval $[\underline{x}, \bar{x}]$, we have: $-\phi'(0) = f(\theta) = 1/(\bar{x} - \underline{x})$, hence

$$\gamma^* = \frac{\bar{x} - \underline{x}}{n}.$$

Comments:

1. We assume that bidders look for the optimal bidding rule among a *limited set* of bidding rules of the form $v_i - \gamma$. That bidders use rules of the form $v_i - \gamma$ could be motivated even without restrictions, by assuming that the common component α is drawn from a diffuse prior.[10,11] Our perspective however is not to argue in favor of a specific shape, on the ground that it is optimal or approximately optimal for some distributions. The forces that shape bid functions are likely to be driven by considerations that lie outside a specific auction model. Rather than endogenizing *all* aspects of behavior, we take as given a shape (additive shading) and endogenize just one aspect of behavior (the extent of shading), with the hope that most economic insights can be captured in this way.

2. The number of bidders participating in the auction affects bidding in two ways: through the chance of winning ($1/n$), and through the dispersion term ($-1/\phi'(0)$). To evaluate the effect the number of bidders on shading, write $\phi_n(y)$ to indicate the probability that θ_i exceeds $\max_{j \neq i} \theta_j$ by more than y when there are n bidders.[12] We denote by γ^*_n the level of equilibrium shading when there are n bidders present. We have the following proposition:

[10] Alternatively, if the distribution over α is flat on a large interval, then, even if bidders look for the optimal strategy among all possible bid functions, then, except near the boundary of the distribution over v_i, learning v_i is not informative about θ. Formally, let g be the density function of α, assumed to be flat over $[\underline{\alpha}, \overline{\alpha}]$; then for any $v_i \in [\underline{\alpha} + \overline{x}, \overline{\alpha} - \underline{x}]$ we have:

$$f(\theta \mid \alpha + \theta_i = v_i) = \frac{f(\theta)g(v_i - \theta_i)}{\int_{y_i} f(y_i)g(v_i - \theta_i)d\theta_i} = f(\theta).$$

[11] This diffuse prior model, along with idiosyncratic terms drawn from independant and uniform distributions, has also been proposed as a tractable affiliated value model in Klemperer (1999, Appendix D)

[12] We assume that $y + \frac{\phi_n(y)}{\phi'_n(y)}$ is increasing in y for all n.

Proposition 2. *Assume $g(\theta) = f$ is centered, symmetric around 0, and single peaked. Define $\beta_n = -\phi'_n(0)$. Then $\gamma^*_n = \frac{1}{n\beta_n}$* and β_n *is a decreasing sequence. The sequence β_n may decrease to 0 if $f(\bar{x}) = 0$. However, as n increases, $n\beta_n$ increases without bound.*

Intuitively, when the number of bidders increases, the winning bidder tends to have a higher realization of θ_i. Since f is single-peaked, the distribution of other bidders' valuations tends to be more dispersed on average, and β_n decreases.

3.2 Existence

In standard auctions, existence of an equilibrium with monotone strategies is a difficult issue in general, and providing economic insights as to when existence fails may be difficult. Our approach deals with shading levels rather than shading functions, and interpretation is easier. As with a standard problem of *price competition with differentiated products*, existence depends on the shape of the "demand" function ϕ.

"Local" deviations are taken care of by first order conditions, and there are two types of "large" deviations that may create difficulties: either shading by a much larger amount, with the hope that the chance to win does not vanish; Or shading by a much smaller amount with hope that the chance of winning will be much greater.

A classic condition that guarantees existence is that ϕ is log-concave.[13] Another condition, weaker but still sufficient, is that $y + \frac{\phi}{\phi'}$ is non decreasing.

$A2$: ϕ is log-concave.

$A2'$: $y + \frac{\phi}{\phi'}$ is increasing.

Proposition 3: *Under either $A2$ or $A2'$, existence of a pure strategy equilibrium is guaranteed.*

Proof: $A2$ implies $A2'$. $A2'$ implies that the best response is uniquely defined, hence the shading level γ^* derived from first order conditions is an equilibrium. QED

Intuitively, one expects ϕ to be S-shaped, so convex on some range. Convexities are potentially problematic because they may induce incentives for large deviations, either upward or downward. A2 and A2' are conditions that limit the extent to which ϕ is convex, making it sufficient to check for first-order conditions.

A typical case where these conditions (and existence) fail is when the density f simultaneously exhibits *some concentration* and *fat tails.* Concentration implies a strong Bertrand competition effect, hence little shading (and little profit) in any tentative equilibrium, while fat tails imply that the chances to win remain nonnegligible even when shading substantially. Thus there exists a force toward large shading: you can take a chance on a large benefit, even if it is at the risk of having little chance of winning.

[13] This means that $Log\phi$ is concave.

Similarly, these conditions will fail when there is significant uncertainty about the dispersion of idiosyncratic terms, as we now illustrate.

Example. Consider an auction with two bidders where idiosyncratic terms are either drawn from a distribution with density g_1 (say, this is state $k = 1$) or g_2 (under state 2), and assume that state $k = 1$ has probability q. Define

$$\phi_k(y) = \Pr{}_k\{\theta_i - \max_{j \neq i} \theta_j > y\}$$

We have:

> **Proposition 4:** *If* $\phi_1'(0) \max_y y\phi_2(y) > \frac{1}{4q(1-q)}$ *then existence of pure strategy equilibria fails.*

In words, existence fails for example when the function $\phi = E\phi_k$ simultaneously exhibit concentration (due to high $\phi_1'(0)$) and large dispersion (high $\max_y y\phi_2(y)$)

Proof: The first order condition implies that a tentative equilibrium shading γ^* that satisfies:

$$\gamma^* = -\frac{q\phi_1(0) + (1-q)\phi_2(0)}{q\phi_1'(0) + (1-q)\phi_2'(0)} \leq \frac{1}{-2q\phi_1'(0)},$$

hence a profit at most equal to $\frac{1}{-4q\phi_1'(0)}$. By picking a large shading level γ that maximizes $y\phi_2(y)$, a player can secure at least $(1-q)\max y\phi_2(y)$. So existence fails when the condition of the Proposition holds. QED

Note that although pure strategy equilibria may fail to exist, relatively simple equilibria in mixed strategies can be constructed. Our working paper (Compte and Postlewaite, 2010) provides an illustration, with bidders randomizing between only two levels of shading.

3.3 Revenue Rankings

We compare below two auction formats: first price and second price auctions. The usual insight concerning private value auctions is that if valuations are affiliated, then second price auction generates more revenue. As for existence our approach simplifies the analysis, and it proposes an alternative interpretation for the comparison.

In a second price auction, the winner, say player i, gets $y = \theta_i - \max_j \theta_j$ in events where y is non negative. Since y is distributed according to the density $-\phi'(y)$ (by definition of ϕ), a bidder's expected gain, which we denote G^{II}, is therefore:

$$G^{II} = \int_{y \geq 0} -y\phi'(y)dy = \int_{y \geq 0} \phi(y)dy.$$

In a first price auction, a bidder's expected gain, which we denote G^I, is equal to

$$G^I = \gamma^*\phi(0) = \frac{[\phi(0)]^2}{-\phi'(0)}.$$

Since the allocation does not change across formats, the seller's revenue is highest when the bidder's expected gain is smallest. So we conclude:

Proposition 5: *Assume $A2'$ holds. The first price auction generates more revenue than the second price if and only if*

$$| \phi'(0) | \int_{y \geq 0} \phi(y)dy > [\phi(0)]^2$$

The first price is thus preferable when the "demand" function combines high concentration (i.e. high $| \phi'(0) |$, implying strong Bertrand competition effect, hence low gains in the first price auction) and substantial dispersion (implying high rents in the second price auction).[14]

Intuitively, high dispersion does not generate high rents for bidders in the first price auction because bidders who happen to get high realization do not know/realize it, so they cannot tailor shading to that event and realize that increased shading would be profitable.

3.4 Releasing information about dispersion

To conclude this Section, we compare the case where bidders have access to information on the dispersion of valuations, to the case where they don't. We model this information as a signal $k \in K$ that may be released (or not) to bidders, signal k having probability q_k. We assume that this information preserves symmetry and define:

$$\phi_k(y) \equiv \Pr(\theta_i - \max_{j \neq i, j \in I} \theta_j > y \mid k)$$

We assume that $A2'$ holds for each ϕ_k and for $\phi = \sum_k q_k \phi_k$. Let γ^* denote the equilibrium shading level when participants do not know k, and γ_k^* the equilibrium shading when they know k. We have:

Proposition 6: $\gamma^* < E\gamma_k^*$

That is, bidders bid more aggressively on average when they do not know k. Hence for the seller, the policy of not releasing information about k generates more revenue.

Proof: From Proposition 1, we have:

$$\gamma_k^* = \frac{\phi_k(0)}{-\phi_k'(0)} \text{ and } \gamma^* = \frac{\sum_k q_k \phi_k(0)}{-\sum_k q_k \phi_k'(0)}. \tag{1}$$

Let $p(k \mid i)$ denote the joint probability that the state is k given that i wins: $p(k \mid i) = q_k \phi_k(0) / \sum_k q_k \phi_k(0)$. We have:

$$\frac{1}{\gamma^*} = \sum_k p(k \mid i) \frac{1}{\gamma_k^*} > \frac{1}{\sum_k p(k \mid i) \gamma_k^*}, \tag{2}$$

[14] Technically, if we let $\Psi(y) = \int_y \phi(x)dx$, the condition is equivalent to $\log \Psi$ being concave around 0, hence it is weaker than both $A2'$ and $A2$.

where the inequality follows from $y \longmapsto 1/y$ being a convex function. Given the symmetry assumption, $\phi_k(0) = 1/n$ for all k so $p(k \mid i) = q_k$, so we conclude. QED

This Proposition immediately extends to the case where the *number of participants is stochastic* and where k provides the number or the identity of the participants, to the extent that symmetry is preserved: all players have equal chance of being a participant.

> **Proposition 7**: *Assume that the set of participants I is stochastic and that k reveals the number of participants $\mid I \mid$. If $\Pr\{i \in I \mid k\}$ is independent of i (hence equal to $\frac{|I|}{n}$), then Proposition 6 holds.*

Proposition 7 thus confirms the results of Matthews (1987) and McAfee and McMillan (1987), and points out that releasing information about the number of participants is analogous to releasing information about dispersion.[15]

Proof: Redefine $\phi_k(y)$ as the ex ante probability of being a participant and winning when the state is k and while shading more than all other participants by exactly y:

$$\phi_k(y) = \Pr\{i \in I \mid k\} \Pr(\theta_i - \max_{j \neq i, j \in I} \theta_j > y \mid k)$$

Equations and inequality (1) and (2) hold unchanged. So does the equality $p(k \mid i) = q_k$ since $\phi_k(0) = \frac{|I|}{n} \frac{1}{|I|} = 1/n$. QED

4 Incorporating information about rank

We motivated our model by arguing that in many contexts, a bidder's own valuation is a poor tool for estimating how his value compares to others' valuations. We do not mean to suggest, however, that bidders cannot form predictions about rank stemming from signals about others that they might receive.

More generally, there are other signals beyond valuations that bidders might have access to and able to process. We illustrate below how such signals can be incorporated in our basic model. We then examine two different signal structures in which bidders receive one of two possible signals, indicating "high rank" or not. We restrict attention to the case of two bidders and show that information about rank can either be pro-competitive (Proposition 8) or anti-competitive (Proposition 9).[16] A pro-competitive effect obtains when the "high rank" signal

[15] Note that the statement assumes that existence of a pure strategy equilibrium obtains, whether k is observed or not. With a stochastic number of bidders, existence of a pure strategy equilibrium for each realization of n with n known to bidders does not guarantee existence of a pure strategy equilibrium in the uncertain case. The reason is identical to that provided in Proposition 4.

This comment that existence may fail with a stochastic number of bidders actually applies to the standard approach as well. We make this precise in the working paper version.

[16] The insight that additional signals may be anti-competitive (and decrease revenues) has been documented by Fang and Morris (2006), in an independent private value auction where

is more likely when rank is indeed higher (i.e. $v_i > v_j$). An anti-competitive effect obtains when the "high rank" signal is delivered if and only if value is *substantially* higher (i.e. $v_i > v_j + \Delta$).

Formally, bidder i's data is defined as:

$$z_i = (v_i, k_i),$$

where $k \in K_i$ and K_i is a finite set. A plausible rule for bidder i, denoted r_γ, will now consist of a vector of shading levels, one for each possible signal k_i. With $\gamma = (\gamma^k)_{k \in K_i}$, we define

$$r_\gamma(v, k) = v - \gamma^k,$$

and assume that the set of rules R consists of all such rules. Equilibrium definition is unchanged.[17]

Consider now the case of two bidders who receive private (and possibly noisy) information about their rank.[18] Specifically, assume two possible signals, i.e. $K_i = \{0, 1\}$, and the following technology, whereby for any valuation vector $v = (v_i)_i$,

$$\begin{aligned} \Pr\{k_i = 1 \mid v_i > v_j\} &= p \text{ and} \\ \Pr\{k_i = 1 \mid v_i < v_j\} &= 1 - p \end{aligned}$$

So when $p = 1/2$, the signal is uninformative, while for $p = 1$, it is perfectly informative of whether i has the higher valuation.

Define (γ_0^*, γ_1^*) as the equilibrium shading levels for each $k_i = 0, 1$. We refer to γ^* as the equilibrium shading level when no signal about rank is available. We have

Proposition 8: Under $A2$, for any $p > 1/2$, we have $\gamma_0^* < \gamma_1^* < \gamma^*$.

The proof is in the Appendix. The intuition is as follows. The private signal creates an asymmetry. A bidder who receives good news might be willing to exploit that signal to bid less aggressively. Given the particular signal structure

valuation may take two values. The insight that additional signals may be pro-competive (and increase revenues) has been documented by Lansberger and al. (2001) and Fang and Morris (2006), in a standard independent private value model with a continuum of valuations, and two signals correlated with rank.

Propositions 8 and 9 thus confirms these insights, in a setting where bidders do not use valuation realizations to make inferences about rank. The propositions also provide some insight about how the nature of the signal structure affects translates into more or less competition.

[17] Such signals could also be introduced in the standard model. The technical difficulty is that when signals are private, bidders have a two dimensional type, and equilibria are then difficult to characterize.

[18] The case where bidders receive public and perfect information about rank has been studied by Landsberger et al. (2001). We extend their insights to the case of imperfect private information about rank, but our main motivation here is to illustrate the simplicity of the approach.

assumed however, it turns out that this is not the case. If $p = 1$ for example, the good news $k = 1$ just moves the winning probability $\phi(y)$ to $2\phi(y)$, thus not altering incentives to shade. For the bidder receiving bad news however, incentives are changed and he wishes to be more aggressive ($\gamma_0^* < \gamma^*$). Under A2, best responses are monotonic, that is, more aggressive behavior from one's opponent triggers a more aggressive response, implying that $\gamma_1^* < \gamma^*$.

Proposition 8 shows the pro-competitive effect of obtaining information about rank. The special signal structure assumed is important however. If a bidder only gets good news when he is far ahead of the other, then in these events, he will be inclined to be less aggressive, and over all, such signal structure may be decrease competition. The next Proposition confirms that intuition.

Formally, fix $\Delta > 0$ and consider the following technology, whereby for any valuation vector $v = (v_i)_i$, bidder i observes $k_i = 1$ if and only if $v_i > v_j + \Delta$, and $k_i = 0$ otherwise. We consider the non-degenerate case where $\phi(\Delta) > 0$.[19] We have:

> **Proposition 9:** Under *A*2, if $\gamma^* < \Delta$ then $\gamma_1^* = \Delta$ and $\gamma_0^* = \gamma^*(1 - 2\phi(\Delta))$. If $\gamma^* < \Delta/2$, then bidders get more expected gains than if they had not received signals.

In words, the bidder who gets good news is less aggressive, the bidder who gets bad news is more aggressive. Overall however, if γ^* is small enough, the latter effect has less impact on expected gains, and on average, being able to observe these signals generates more expected gains for bidders.

Discussion

A standard division in studying auctions is whether values are *independent or correlated*. We argue below that this division may be useful to analysts, but from an agent's perspective, a more useful dividing line may be *whether or not he gets and exploits rank related signals.*

We have represented values as the sum of two *random* components, a common component α and an idiosyncratic component θ_i. Given this representation, the classic *independent private value* environment corresponds to the case where bidders observe both the value v_i and the common component α, that is,

$$z_i = (v_i, \alpha).$$

The case where agents lack precise knowledge of the common component α corresponds to the standard *correlated* private value environment. That lack of knowledge is generally modelled by assuming α is a random variable and that the agent is uncertain about its realization. This uncertainty can be modelled as each bidder receiving a noisy estimate of α, say

$$\beta_i = \alpha + \xi_i.$$

[19] This ensures that the events $k_i = 1$ arise with positive probability

A player's data would then consist of

$$z_i = (v_i, \beta_i),$$

with the independent private value case being the degenerate case with no noise.

From a purely *technical perspective*, the independent private value model is relatively simple: the analyst knows the complete structure of the model, and he finds it easy to aggregate the two pieces of information v_i and α into bid functions of the form $r(v_i, \alpha) = b(v_i - \alpha)$, and among those, to derive an optimal bid function $r^*(v_i, \alpha)$.

The correlated case is more complex because how exactly one ought to aggregate these two signals is difficult, even for the analyst that knows all distributions. The practice in general consists in ignoring this difficulty by assuming that no such estimate β_i is observed.[20] As for independent values case, however, each realization v_i comes with *implicit* information about rank and dispersion of other's values (and information), and the agent ends up being able to exploit that information, or behaving as if he could.

Alternatively, let us take *a bidder's perspective.* First, even if interested in his rank, he may find impossible to come up with some relevant or reasonably accurate reference point (α) to which his value v_i could be compared. Second, even if he could form an estimate (β_i) of such reference point, this estimate is likely to be noisy. How exactly he ought to combine these two signals v_i and β_i seems extraordinarily complex, as even an analyst knowing the distributions cannot figure it out.[21]

Our perspective is to propose models that reflect the cognitive difficulties that agents face, starting from mildly sophisticated ones and then possibly investigating what further sophistication brings. The simplest environment examined has been that in which bidders don't form, nor try to use, nor end up using, rank-related observations (Section 2 and 3). A more elaborate environment is one where bidders get and attempt to use rank-related information in a crude way (only two signals), and this Section has illustrated different ways in which such information may alter the strategic interaction, either decreasing or promoting competition.

The standard independent private value model pushes sophistication even further, as agents end up using rank and dispersion related signals in very subtle ways, and one might question what this more elaborate model adds to the insights of this Section.

5 Noisy Estimates

Throughout this section, we drop the private value assumption ($z_i = v_i$) and move back to our basic model, assuming that z_i is a noisy estimate of v_i. We

[20] There are exceptions: Fang and Morris (2006) for example.

[21] A crude but plausible rmental processing rule could be that he thinks his rank is high ($k = 1$) or not that high ($k = 0$) depending on whether $v_i - \beta_i$ is higher than some threshold h, and thus end up with "data" $z_i = (v_i, k)$.

assume that the noise terms ε_i are centered and drawn from independent and identical distributions. We look for a symmetric equilibrium of the first price auction. Denote by $\phi_\varepsilon(y)$ the probability that z_i exceeds $\max_{j\neq i} z_j$ by more than y, that is,

$$\phi_\varepsilon(y) = \Pr(z_i - \max_{j\neq i} z_j > y).$$

Denote by $v_i(\gamma_i, \gamma)$ the expected payoff that bidder i derives when he shades by γ_i while others shade by γ. Bidder i obtains a payoff equal to $v_i - (v_i + \varepsilon_i - \gamma_i)$ when he wins, so we have:

$$v(\gamma_i, \gamma) = (\gamma_i - E[\varepsilon_i \mid z_i - \max_{j\neq i} z_j > \gamma_i - \gamma])\phi_\varepsilon(\gamma_i - \gamma).$$

Define

$$\psi_\varepsilon(y) = E[\varepsilon_i \mid z_i - \max_{j\neq i} z_j = y].$$

First order conditions immediately yield:

Proposition 10: *In a symmetric equilibrium, bidders shade their bid by*

$$\gamma_\varepsilon^* = \frac{1}{-n\phi'_\varepsilon(0)} + \psi_\varepsilon(0).$$

Shading thus has two components. The first term is analogous to that derived in the private value case: $(\frac{1}{n})$ corresponds to the expected probability of winning, and $\frac{1}{-\phi'_\varepsilon(0)}$ captures how bidders take advantage of the dispersion in valuations. The second term $\psi_\varepsilon(0)$ takes into account the fact that more optimistic bidders tend to win the auction (i.e., the winner's curse), and rational bidders should correct for that.

5.1 Comparative statics.

When estimates are noisier, dispersion of estimates increases and bidders can take advantage of that by shading their bids more. We provide below examples where shading increases. Increased shading however does not necessarily translate into more gains for bidders, because a bidder only gains $\gamma_\varepsilon^* - \varepsilon_i$ in the event he wins. However, we also illustrate below that bidders can be *unambiguously* better off when estimates are noisier. To fix ideas we assume that each θ_i is uniform on an interval of size Δ.

Example 1: Consider the case of two bidders. By symmetry $\psi_\varepsilon(0) = 0$, and since without noise, $-\phi'(y)$ is maximum at 0, $-\phi'_\varepsilon(0)$ may only be smaller. It follows that $\gamma_\varepsilon^* > \gamma^*$, that is, *shading increases with noise.*[22] The effect on expected gains is ambiguous in general. However we have:

[22] This argument does not require uniform distributions. It holds if the θ_i's are drawn from centered and single peaked distributions.

Proposition 11: Assume noise term has uniform distribution. Then *for Δ small enough, bidders are better off with noise than without.*

Example 2: With more than 2 bidders, and a simple noise structure, we illustrate how pressure towards higher shading obtains. We assume ε can take two values, $\bar{\varepsilon}$ (with probability p) or $\underline{\varepsilon}$, so that bidders are either optimistic or pessimistic. Let $\Delta_\varepsilon = \bar{\varepsilon} - \underline{\varepsilon}$ and assume that $\Delta_\varepsilon > \Delta$. In a symmetric equilibrium, bidder i may only win when he is optimistic ($\varepsilon_i = \bar{\varepsilon}$) or when all bidders are pessimistic ($\varepsilon_j = \underline{\varepsilon}$ for all j). Define $\widetilde{n}$ as the random variable that gives the number of bidders who have a chance to win for each realization $(\varepsilon_i)_i$, that is:

$$\widetilde{n} = \#\{i, \varepsilon_i = \max_j \varepsilon_j\}$$

We have:

Proposition 12: *In a symmetric equilibrium, we have:*

$$\gamma_\varepsilon^* \geq E[\varepsilon \mid \varepsilon_i = \max_j \varepsilon_j] + \Delta E \frac{1}{\widetilde{n}}. \qquad (3)$$

Intuitively, the first term corresponds to the expected "optimism" of the bidder conditional on winning, and that term gets close to $\bar{\varepsilon}$ when n increases. The second term describes how bidders further shade their bids by exploiting the dispersion of estimates. Compared to the case without noise where they would shade by Δ/n, bidders shade more because they are facing less intense competition: because only optimistic bidders may win (except in the event all are pessimistic), a bidder is endogenously facing fewer competitors.[23]

Example 3. We conclude with the two bidder case, explaining why noise may actually benefit the seller (for some distributions). Intuitively, noise implies more dispersed bids, and the extent to which bidders exploit that dispersion depends on $\phi_\varepsilon'(0)$, a local condition. Now the benefits of dispersion depends on the whole shape of ϕ_ε. If $\theta_i - \theta_j$ has a density h which is relatively flat around 0 (i.e. $h''(0) << h(0)$), then adding a small noise has a second order effect on shading. However, it has a first order effect (and positive effect) on the bid of the winner. Hence the seller's revenue rises

5.2 Comparison with second price

We now compare with the second price auction. We start by deriving equilibrium shading. Next we show that the previous insight, noisier estimates may induce higher rents, also holds with second price auctions.

[23] Letting $\rho = (1-p)^n/p$, it can be shown that

$$E[\varepsilon \mid \varepsilon_i = \max_j \varepsilon_j] = \bar{\varepsilon} - \frac{\rho}{1+\rho}(\bar{\varepsilon} - \underline{\varepsilon}) \text{ and } E\frac{1}{\widetilde{n}} = \frac{1}{pn(1+\rho)}.$$

In a second price auction, bidder i now gains $v_i - (\max z_j - \gamma) = z_i - \max z_j + \gamma - \varepsilon_i$ when he wins. Letting $h_\varepsilon(y) = -\phi'_\varepsilon(y)$, we have

$$v(\gamma_i, \gamma) = \int_{y \geq \gamma_i - \gamma} y h_\varepsilon(y) dy + (\gamma - E[\varepsilon_i \mid z_i - \max_{j \neq i} z_j > \gamma_i - \gamma])\phi_\varepsilon(\gamma_i - \gamma).$$

First order conditions now imply:

Proposition 13: *In the second price auction, in a symmetric equilibrium, bidders shade by*

$$\gamma_\varepsilon^{SP} = \psi_\varepsilon(0).$$

Bidders thus correct for winner's curse effect in the same way. Going from second to first price, the change in bidding thus only stems from bidders taking advantage of the dispersion in estimates.

As with the first price auction, noisier estimates may help bidders. Conditional on winning, a bidder obtains:

$$G^{II} = E[\theta_i - \max z_j \mid z_i > \max z_j] + \gamma_\varepsilon^{SP}.$$

With two bidders, $\psi_\varepsilon(0) = 0$ by symmetry, so $\gamma_\varepsilon^{SP} = 0$, and we have:

$$G^{II} = E[\theta_i - \theta_j \mid z_i > z_j] - E[\varepsilon_j \mid z_i > z_j].$$

The first term is the efficiency gain that the winner brings. The second term is an additional rent that the winner obtains because the loser is pessimistic on average. As noise increases, the first term decreases because the allocation is less efficient. The second term however increases. When the idiosyncratic components θ_i get more concentrated, the second term prevails, and the agents get rents that they would not get in the private value case.

The same is true with many bidders. As the idiosyncratic components θ_i get more concentrated, the winner's gain tends to

$$G_0^{II} = E[\max \varepsilon_j \mid \max_j \varepsilon_j = \varepsilon_i] - E[\max \varepsilon_j \mid \max_{j \neq i} \varepsilon_j < \varepsilon_i],$$

which is strictly positive. So dispersion of estimates always increases rents when idiosyncratic components are concentrated. These rents result from the dispersion of estimates.

We conclude with the following proposition:

Proposition 14: *Assume two bidders, with θ_i and ε_i drawn from centered distributions. Then noisy estimates never benefit the seller.*

Intuitively, noise can sometimes improve the seller's revenue, because both bidders may be optimistic. However, the loser determines the price and on average the loser is more pessimistic.

Proof: With two bidders, bidders bid their estimate z_i, so the seller obtains $G = E\alpha + E[z_j \mid z_j < z_i]$, so since θ_i and ε_i are centered (around 0), we have $G = E\alpha - \frac{1}{2}E[z_i - z_j \mid z_i > z_j]$. We show below that noise inflates $H \equiv E[z_i - z_j \mid z_i > z_j]$, thereby concluding the proof. We have

$$H = -\int_{y>0} y\phi'_\varepsilon(y)dy = \int_{y>0} \phi_\varepsilon(y)dy = \int_{y>0}\int_\varepsilon \phi(y+\varepsilon)h(\varepsilon)d\varepsilon dy$$

where h denotes the density over $\varepsilon = \varepsilon_i - \varepsilon_j$. Define $\varphi(\varepsilon) = \int_{y>0} \phi(y+\varepsilon)dy$. φ is convex, implying that $H = \int_\varepsilon \varphi(\varepsilon)h(\varepsilon)d\varepsilon > \varphi(0) = \int_{y>0} \phi(y)dy$, as desired. **QED**

5.3 Discussion

Dispersion rents or Information rents?

We often refer to information rents to describe the gains that an "informed" player gets. A more appropriate qualification might be *"dispersion rents"*: it is not information that generates rents, but bid dispersion, and in our model, noisier estimates translate into higher bid dispersion.

In a standard auction model, the opposite would be true, as noisier estimates translate into **less** dispersed posteriors (by a regression to the mean effect), and therefore greater competition when symmetry is assumed. The latter conclusion however is (in our view) an unfortunate artifact of the model, and of the implicit assumption that agents know (or behave as if they knew) all distributions: as noise increases, value estimates decrease in importance and more weight is put on priors.

Common values

In modelling auctions, the distinction between private and common values is often seen as a key dividing line. In common or interdependent value auctions, the bids of others reveal information about one's own valuation, and a bidder ought to care about those inferences. The omniscient bidder will indeed find that warning useful, and find the appropriate bid function. To most bidders however, the precise ways in which preferences are interdependent are probably obscure, and making sense of this warning is likely out of reach.

From a less sophisticated bidder's perspective, a more useful dividing line may be whether he is subject to estimation errors or not. If he is subject to estimation errors, he should exert caution because he is subject to a selection bias: he is more likely to win when the error is positive.[24] That warning is not specific to auctions. It arises for any decision problem in which he compares an alternative that is easy to evaluate (not buying) to one that is more difficult to evaluate (Compte and Postlewaite 2012).

Now the level of caution depends on context, and indeed, the degree of interdependence then matters. If idiosyncratic components are less dispersed

[24] The view that the winner's curse results from a selection bias appears in Capen and al. (1971). Compte (2001) examines the effect of increasing the number of bidders on this selection bias, in the context of the second price auction.

(which can be interpreted as values being more interdependent), the estimation errors carry more weight and caution has to increase.[25]

6 Discussion and Related literature

6.1 Simple rules and strategy restrictions.

Our approach studies auctions where bidders only compare simple rules. Following Rothkopf (1969), a number of authors have analyzed auctions where bidders use multiplicative bidding strategies. More recently, Klemperer (1999), Compte (2001), Compte and Postlewaite (2010) have considered the additive structure proposed here.

Auction models proposing simple bidding rules generally view them as arising from standard equilibrium considerations. These models put a special structure on the joint distributions over valuations ensuring that using simple rules is optimal *across all possible rules* (given that others use simple rules as well). This special structure typically assumes that valuations are functions of a common component drawn from a *diffuse prior.* We have proposed a different perspective, based on *direct restrictions* on the strategy space. While both perspectives capture the notion that bidders might have difficulties extracting rank information from their valuations, we propose the path that restricts the strategy space for various reasons:
(i) Our perspective is not to argue in favor of a specific shape, on the ground that it is optimal or approximately optimal for some distributions. The forces that shape bid functions are likely to be driven by considerations that lie outside a specific auction model. Rather than endogenizing *all* aspects of behavior, we take as given a shape (additive shading) and endogenize just one aspect of behavior(the extent of shading), with the hope most economic insights can be captured in this way.
(ii) We believe that strategy restrictions are a useful tool not only in auctions but also in other strategic environments – dynamic ones for example – for which the diffuse prior assumption would have no equivalent;
(iii) From a positive perspective, we are implicitly attempting to deal with agents having *limited knowledge* of the environment they are facing (say the joint distribution over valuations). The standard way to deal with limited knowledge would be to follow the traditional Bayesian route that takes anything that is not known as a random variable, and then possibly assume that agents receive signals correlated with the realized joint distribution. However, following Heiner (1983), one would expect that limited knowledge be synonymous to *lesser* sophistication, *not more.* Restrictions on the strategy space capture such a bound on sophistication that limited knowledge would seem to call for.

[25] Caution also has to increase when the number of bidders increases (Compte 2001).

6.2 Existence issues and revenue rankings

Existence. In standard auction models, the concern has been to establish existence of equilibria in monotonic strategies. The main difficulty is that based on his signal/valuation, the omniscient bidder makes sophisticated inferences about the distribution over other bidders' valuations, and it is not guaranteed that his best response will be monotonic in his signal. Affiliation is a condition that ensures that best responses are indeed monotonic in one's signal (when other players use monotonic strategies).[26]

Our analysis highlights that existence issues arise even when players are much less sophisticated (in particular monotonicity of bid functions is not an issue – monotonicity is assumed). The reasons are analogous to those that prevent existence of pure strategy equilibria in models of price competition with differentiated products (Caplin and Nalebuff 1986). The basic problem is that a bidder may simultaneously have strong incentives to be slightly more aggressive than others (buying a substantial chance of winning at little cost), and strong incentives to take a chance and opt for large shading, betting that the other players have a substantially lower valuation. Whether both strategies turn out to be attractive depends on the shape of the "demand function" ϕ that characterizes the dispersion of idiosyncratic components.

Revenue rankings. Revenue equivalence between first and second price auctions holds for independent private value auctions (Myerson, 1981). Second price generated more revenue when private valuations are affiliated (Milgrom Weber, 1982). Economic intuition for these comparisons is somewhat difficult to provide. Milgrom Weber (1982) appeal to a "linkage principle", which itself seems to apply as soon as valuations are positively correlated. However affiliation is stronger than positive correlation and De Castro (2007) reports simulations where valuations are assumed to be positively correlated and yet first price auction generate greater revenue.

Standard revenue comparisons are made assuming omniscient bidders who can finely extract rank and dispersion information from their valuations and adjust bidding accordingly. Our analysis points out that these comparisons depends on bidders being omniscient, and that one might want to understand what drives revenue comparisons when agents are less sophisticated.[27]

In this vein, Proposition 5 provides a simple characterization based on the shape of the demand function ϕ. Proposition 5 also helps build intuition. In essence, first price is better when idiosyncratic elements have a small chance (but not too small) of being sufficiently spread out. The seller benefits from this situation because "high valuation" bidders cannot take advantage of that dispersion: they cannot extract information about rank and dispersion from their valuation, so effectively, they do not know that they are "high valuation" bidders (if they could they would shade more, thereby decreasing sellers' rev-

[26] See for exampe Athey (2001), Lizzeri and Persico (1995), and Reny and Zamir (2004)

[27] Fang and Morris (2006) also question the revenue rankings; by providing agents with an additional private signal correlated with the other player's valuation, symmetry is broken, and inefficiencies naturally arise in the first price auction.

enue).

Finally, Proposition 5 also sheds light on revenue ranking in *standard models*: if the common component is drawn from a flat distribution with very large support (or a diffuse prior), then even omniscient bidders cannot extract information about rank, and the intuition above applies.

Comparative statics.

In standard models, the seller's revenue depends on the fine inferences about own valuation, rank and dispersion that agents make based on their value estimate. Access to other signals or information beyond one's value estimate modifies these inferences, hence also the seller's revenue. When the additional information released is affiliated with the bidders' signals, the effect on revenue is nonnegative (Milgrom Weber, 1982).

We explored three kinds of comparative statics that separate the effect of further information on dispersion, on rank, and on own valuation. Propositions 6 and 7 show that the policy of disclosing information about the dispersion of valuations (or the number of bidders) while keeping symmetry has an anti-competitive effect, beneficial to bidders (hence detrimental to seller's revenue – since the allocation does not change).

Propositions 8 and 9 show that providing information about rank to bidders may have an ambiguous effect on bidders' gains and seller's revenue. In particular, revenue may decrease despite the fact that information about rank is positively correlated with valuation (which thus provides another illustration of the difference between affiliation and positive correlation).

Finally, we have examined the effects of the quality of estimates on bidders gains and seller's revenue. When estimates are noisy, shading in the first price auction is driven two factors: the dispersion of estimates, and the selection bias (i.e. the error conditional on winning by a zero margin). Propositions 11 and 12 illustrate the two channels by which noisier estimates can increase shading: a weaker Bertrand competition effect (due to higher dispersion of estimate), and a stronger selection bias. These two effects may result in bidders benefiting from noise, hence since efficiency is reduced, smaller revenue for sellers. Example 3 however shows that depending on the shape of the distribution over idiosyncratic elements, sellers' revenue may actually increase when estimates are noisier. Nevertheless, and contrasting with Ganuza 2004 who shows that with two bidders, ignorance promotes competition, we find that with two bidders, noisier estimates always hurt the seller in the second price auction. The reason for this difference has been explained earlier: standard models implicitly assume that weaker information translates into more concentrated posteriors, hence stronger competition.

7 Conclusion

We have proposed a model in which players only consider a limited set of strategies. This limitation can be interpreted as a bound on rationality, or a bound on the ability to determine what strategies are optimal when the strategy set is large. It can also be viewed as an analyst's device or methodological tool to deal with agents' lack of detailed knowledge over prior distributions, or to deal with agents' inability to use effectively this prior information.

We see various benefits from the approach. Starting from the sophisticated end of the spectrum, it offers a way to check the robustness of insights derived from standard models. It may constitute a useful alternative to the robustness literature, in particular when there are no easy or tractable ways to enrich the types space. It may also shed a different light on known results.

Starting from the other end of the spectrum (in terms of sophistication), it offers a more parsimonious theory of auctions, that can always be amended by increasing sophistication, but up to a degree that the analyst considers plausible.

Finally, it questions what poor information means. We have taken the view that poorer information meant larger discrepancy between value and value estimate. Whether poorer information leads to more dispersed estimates or opinions is probably a matter of context, or at least an empirical question. The existence of a natural reference point is probably a pre-requisite to obtain that poor information leads to lesser dispersion. Given the well documented inability for agents to correctly take into account priors in forming beliefs (Tversky and Kahneman, 1974), the existence of such reference point may not be sufficient.

References

Athey, S., 2001, Single Crossing Properties and the Existence of Pure Strategy Equilibria in Games of Incomplete Information, *Econometrica* 69 (4): 861-890.

Bergemann, D and S. Morris, 2005, Robust Mechanism Design. *Econometrica*, 73,1771-1813.

Bergemann, D., M. Pesendorfer (2007), Information structures in optimal auctions, *Journal of Economic Theory* 137, 580-609

Block, H. D. and J. Marschak, 1960, "Random orderings and stochastic theories of responses", in I. Olkin et al., (Eds.), Stanford University Press, Stanford, pp. 97-132.

Capen, E.C. , R.B. Clapp, and W. M. Campbell, 1971, Competitive Bidding in High Risk Situations, *Journal of Petroleum Technology*, 23, pp 641-53.

Caplin, A. and B. Nalebuff 1986, Multi-Dimensional Product Differentiation and Price Competition, Oxford Economic Papers , New Series, Vol. 38, Supplement: Strategic Behaviour and Industrial Competition, pp. 129-145

de Castro, L. (2007) Affiliation, Equilibrium Existence and the Revenue Ranking of Auctions, mimeo

Chatterjee, K. and W. Samuelson (1983) "Bargaining under incomplete information," *Operations Research* 31, pp. 835 -851.

Compte, O, 2001, The winner's curse with independant values, mimeo

Compte, O. and A. Postlewaite, 2010, Auction notes, mimeo

Compte, O and A. Postlewaite (2012) Cautiousness, mimeo

Fang, H. and Morris, S. 2006, Multidimensional private value auctions, *Journal of Economic Theory* 126, 1 – 30

Ganuza, J.-J. (2004), Ignorance Promotes Competition: An Auction Model with Endogenous Private Valuations, *The RAND Journal of Economics*, Vol. 35, No. 3, pp. 583-598

Geanakoplos, J. 1989, Game Theory Without Partitions, and Applications to Speculation and Consensus, Discussion paper n. 914, Cowles Foundation

Jehiel, P. 2005, "Analogy-Based Expectation Equilibrium," *Journal of Economic Theory* 123(2): 81-104.

Kamenica, E. and Gentzkow, M., 2011, Bayesian Persuasion, *American Economic Review*, 101 : 2590-2615

Klemperer, P. 1999 Auctions: A guide to the literature, Journal of Economic Surveys, 13, 3, pages 227–286.

Landsberger, M., Rubinstein, J., Zamir S. and E. Wolfstetter (2001): “First price auctions when the ranking of valuations is common knowledge, *Review of Economic Design* 6, 461-480.

Lindsey, J.H., W. Samuelson and R. Zeckhauser, 1996, Selling procedures with private information and common values, Management Science February 1996 vol. 42 no. 2 220-231

Lizzeri, A. and N. Persico, 1995, Existence and Uniqueness of Equilibrium In First Place Auctions and War of Attrition with Affiliated Values, Discussion Papers 1120, Northwestern University, Center for Mathematical Studies in Economics and Management Science.

E. Maskin and J. Riley (1984) “Optimal auctions with risk averse bidders,” *Econometrica* 6, pp. 1473-1518.

S. Matthews (1987) “Comparing auctions for risk averse buyers: A buyer’s point of view,” *Econometrica* 55, No. 3 , pp. 633-646.

R. P. McAfee and J. McMillan (1987) "Auctions with a stochastic number of bidders," *Journal of Economic Theory*, 43, pp. 1-19.

Milgrom, P. 1981. "Rational Expectations, Information Acquisition, and Competitive Equilibrium," Econometrica 49, pp. 921-43.

Milgrom, P. and R. Weber. 1982. "A Theory of Auctions and Competitive Bidding," Econometrica 50, pp. 443-59.

Morris, S and Shin, H. S. (2006) Global games : theory and applications, Chapter 3, in Cambridge Collections Online

Myerson, R., 1981. "Optimal Auction Design," *Math. Op. Res. 6*, pp. 58-73.

Osborne, M. J. and A. Rubinstein, 1998, "Games with Procedurally Rational Players," *American Economic Review*, vol. 88(4), pages 834-47.

Reny, P. and S. Zamir, 2004, On the Existence of Pure Strategy Monotone Equilibria in Asymmetric First-Price Auctions, *Econometrica* , 72, pp. 1105-1112

Rothkopf, M. H., 1969, A Model of Rational Competitive Bidding, *Management Science* , Vol. 15, No. 7, Theory Series, pp. 362-373

Rubinstein, A., 1989, The Electronic Mail Game: Strategic Behavior Under Almost Common Knowledge, *The American Economic Review* , 79, pp. 385-391

Rubinstein, A. 1991, Comments on the interpretation of game theory, *Econometrica* 4, pp. 909-924.

Tversky, A. and D. Kahneman, 1974, Judgment under Uncertainty: Heuristics and Biases. *Science* 185 (4157): 1124–1131.

Wilson, R. (1987), "Game-Theoretic Analyses of Trading Processes," in: Advances in Economic Theory: Fifth World Congress, Truman Bewley (ed.); Chapter 2, pp. 33-70. Cambridge University Press.

Appendix.

A. Buyer/Seller

We consider a seller with a value v_1 for the object to be sold, and a buyer with a value v_2 for the object. As before, we assume that

$$v_i = \alpha + \theta_i.$$

We are interested in comparing three selling mechanisms: take-it-or-leave it offer by the seller, take-it-or-leave it offer by the buyer and the split-the-difference mechanism (Chatterjee and Samuelson (1983)),[28] under the assumption that players are restricted to bidding rules of the form $r(v) = v + \gamma$, again capturing the idea that players cannot disentangle common and private components. The question asked is thus similar to that addressed by Lindsey et al. 1996.

Define

$$\phi(y) = \Pr\{\theta_2 - \theta_1 \geq y\}$$

and let $S(y)$ denote the expected surplus that results when transactions take place if and only if $\theta_2 > \theta_1 + y$, that is

$$S(y) = \int_{x>y} -x\phi'(x)dz.$$

We have:

> **Proposition 15**: Let $a^* = \arg\max y\phi(y)$. Whether the seller or the buyer makes a take-it-or-leave-it offer, or where players adopt the split-the-difference mechanism, the expected surplus is identical and equal to $S(a^*)$. *The seller prefers to let the buyer make the offer when*
>
> $$\int_{y>a^*} \phi(y)dy > a^*\phi(a^*).$$
>
> Otherwise he prefers to make the offer. Neither player finds the split-the-difference mechanism attractive.

As for the first and second price auction comparison, the critical issue is the extent to which ϕ has a fat tail. If ϕ has a fat tail, it is preferable to let the other player makes the offer: since "high valuation" buyers cannot learn (by assumption) from their valuation that they have indeed a high valuation, they tend to leave high rents to the seller.

B. Uniform and discriminatory auctions

We now consider k objects for sale with n potential buyers, all interested in buying only *one* unit. As before we assume that $v_i = \alpha + \theta_i$, and we are interested in comparing two selling mechanisms: The uniform auction in which

[28] Under the split the difference mechanism (see Chatterjee and Samuelson (1983)), the buyer and seller simultaneously offer respectively prices p_1 and p_2. In the event $p_2 - p_1 \geq 0$, the transaction takes place at price $p = (p_1 + p_2)/2$, otherwise it does not take place.

the seller price is set at the $k+1^{th}$ bid, and the discriminatory auction where bidders pay their bids.

Proposition 5, which ranks revenues from first and second price auctions, easily extends. Denote by $\theta_i^{(k)}$ the k^{th} largest idiosyncratic term among bidders other than i, and define

$$\phi(y) = \Pr\{\theta_i - \theta_i^{(k)} \geq y\}$$

Our analysis of the first price auction extends to the discriminatory auction readily extends to the discriminatory auction with this new definition for ϕ, and we have:

Proposition 16: *Assume A2′ holds for ϕ. The discriminatory auction generates more revenue than the uniform auction if and only if*

$$| \phi'(0) | \int_{y \geq 0} \phi(y) dy > [\phi(0)]^2$$

The proof is identical to that of Proposition 5, and the intuition is the same. High dispersion (fat tail for ϕ) does not generate high rents for bidders in the discriminatory auction because bidders who happen to get high realization do not know/realize it, so they cannot tailor shading to that event and realize that increased shading would be profitable.

C. Proofs

Proof of Proposition 2: We already know that $\beta_n = -\phi'(0)$. By definition, $\phi(y) = \Pr(\theta_i - \max_{j \neq i} \theta_j \geq y) = \int_x f(x)(F(x-y))^{n-1} dx$, thus implying:

$$-\phi'(0) = \int (n-1)(f(x))^2 [F(x)]^{n-2} dx. \tag{4}$$

Integrating by parts the right hand side of (4) and observing that $f'(x) = f'(-x)$ and $F(-x) = 1 - F(x)$, we obtain:

$$\beta_n = f(\overline{\theta}) + \int_{\underline{x}}^{0} f'(x)([1-F(x)]^{n-1} - [F(x)]^{n-1}) dx.$$

Now for any given $p < 1/2$, define $\Delta_n = (1-p)^n - p^n$, and $k_n = \Delta_n/(1-2p)$. We have $k_1 = 1$. Since

$$\Delta_{n+1} = (1-p)\Delta_n + (1-2p)p^n = (1-2p)((1-p)k_n + pp^{n-1})$$

This directly implies $k_n > p^{n-1}$ for all n, which further implies that k_{n+1} is an average of k_n and p^{n-1}, hence k_n is a strictly decreasing sequence. Applying this observation to each $p = F(x)$, we obtain that β_n is also a strictly decreasing sequence. To check the last assertion, assuming for example that $f'(\underline{x}) > 0$, it is sufficient to take a Taylor expansion of $F(x)$ close to $\underline{x}$ and then observe that $\beta_n \geq O(1/n)^{1/2}$. QED

Proof of Proposition 8
Define:

$$\bar{\phi}(y) = \Pr\{\theta_i > \theta_j + y \mid \theta_i > \theta_j\} \text{ and } \underline{\phi}(y) = \Pr\{\theta_i > \theta_j + y \mid \theta_i < \theta_j\}.$$

We have $\bar{\phi}(y) = 2\phi(y)$ for $y \geq 0$, $\bar{\phi}(y) = 1$ for $y \leq 0$, and by symmetry $\underline{\phi}(y) = 1 - \bar{\phi}(-y)$. The event $(k_1, k_2) = (1, 1)$ has probability $p(1-p)$, and conditional on this event, there is equal chance that $\theta_1 > \theta_2$ or $\theta_1 < \theta_2$. The event $(k_1, k_2) = (1, 0)$ has probability $p^2 + (1-p)^2$, and conditional on this event, $\theta_1 > \theta_2$ has probability $p^2/(p^2 + (1-p)^2)$. So for bidder 1, the value from shading by γ in event $k_i = 1$ is:

$$V^1(\gamma, \sigma) = \gamma[p^2\bar{\phi}(\gamma-\gamma^0)+p(1-p)\bar{\phi}(\gamma-\gamma^1)+p(1-p)\underline{\phi}(\gamma-\gamma^1)+(1-p)^2\underline{\phi}(\gamma-\gamma^0)].$$

Similarly, the value from bidding γ in event $k_i = 0$ is:

$$V^0(\gamma, \sigma) = \gamma[p^2\underline{\phi}(\gamma-\gamma^1)+p(1-p)\underline{\phi}(\gamma-\gamma^0)+(1-p)p\bar{\phi}(\gamma-\gamma^0)+(1-p)^2\bar{\phi}(\gamma-\gamma^1))].$$

Defining $z = a^1 - a^0$, the first order conditions become:

$$\begin{aligned} \gamma^1 &= \frac{p^2\phi(z) + p(1-p)\phi(0)}{-[p^2\phi'(z) + p(1-p)\phi'(0)]} \text{ and} \\ \gamma^0 &= \frac{p^2(1-2\phi(z)) + 2(1-p)\phi(0)}{-[2p^2\phi'(z) + 2(1-p)p\phi'(0)]}. \end{aligned}$$

The equilibrium difference $z^* = \gamma^1 - \gamma^0$ thus solves:

$$z^* = \frac{4p^2\phi(z^*) - (p^2 + (1-p)^2)}{-[2p^2\phi'(z^*) + 2(1-p)p\phi'(0)]}.$$

So z^* is strictly positive. Without information about rank, $\gamma^* = -\frac{\phi(0)}{\phi'(0)}$. Since $y + \frac{\phi(y)}{\phi'(y)}$ under $A2'$ (a fortiori under A2), we get $\gamma^0 < \gamma^*$. And since $\frac{\phi(y)}{\phi'(y)}$ is increasing under A2, we get $\gamma^1 < \gamma^*$. QED

Proof of Proposition 9: Let $q \equiv \Pr\{v_i > v_j + \Delta\} = \phi(\Delta)(< 1/2)$. There are three possible events, $k_1 = 1$, $k_2 = 1$ and $k_1 = k_2 = 0$, with respective probabilities q, q and $1 - 2q$. We look for an equilibrium where $\gamma_0^* \leq \gamma^*$. If i observes $k_i = 1$, player j must have seen $k_j = 0$, so i solves $\max_\gamma(\Delta, \frac{\gamma\phi(\gamma-\gamma_0^*)}{q})$. He thus finds optimal to shade by Δ. If player i observes $k_i = 0$ (which means $v_2 - v_1 < \Delta$), he wins only in events where $v_i - v_j \geq \gamma - \gamma_0^*$, so he gains $\frac{\gamma(\phi(\gamma-\gamma_0^*)-q)}{1-q}$. First order conditions yields $\gamma_0^* = \frac{\phi(0)-\phi(\Delta)}{-\phi'(0)}$, hence the desired conclusion. Allocation in unchanged. The winner shades by Δ under events $k_1 = 1$ and $k_2 = 1$, and by γ_0^* under event $k_1 = k_2 = 0$. Expected gains are thus equal to

$$2q\Delta + (1-2q)^2\gamma^*$$

hence when $\gamma^* < \Delta/2$ they are larger than γ^*. QED

Proof of Proposition 10: Define $x = z_i - \max_{j \neq i} z_j$ and let $h(\varepsilon_i, x)$ denote the joint distribution over ε_i and x, and let $H(y) = \int_{x \geq y} \varepsilon_i h(\varepsilon_i, x) d\varepsilon_i dx$. Note that $\psi_\varepsilon(y) = \frac{-H'(y)}{-\phi'_\varepsilon(y)}$. We have:

$$v(\gamma_i, \gamma) = \int_{x \geq \gamma_i - \gamma} (\gamma_i - \varepsilon_i) h(\varepsilon_i, x) d\varepsilon_i dx = \gamma_i \phi_\varepsilon(\gamma_i - \gamma) - H(\gamma_i - \gamma).$$

The first order condition thus yields the desired conclusion. QED

Proof of Proposition 11: Denote Δ_ε the support of the noise term. In the auction without noise, bidders gains tend to 0 when Δ tends to 0. In the auction with noise, bidders get $\gamma^*_\varepsilon - \varepsilon_i$ in the event they win, so their expected gain is:

$$\gamma^*_\varepsilon - E[\varepsilon_i \mid \varepsilon_i + x_i > \varepsilon_j + x_j]$$

By symmetry, $\psi_\varepsilon(0) = 0$, and at the limit where Δ is very small, γ^*_ε tends $\frac{\Delta_\varepsilon}{2}$.[29] The second term tends to $\frac{\Delta_\varepsilon}{6}$, so bidders total payoff tends to $\frac{\Delta_\varepsilon}{3}$, hence remains bounded away from 0. QED

Proof of Proposition 12: Let $\rho = (1-p)^n/p$ and denote by A the event $\{\varepsilon_i = \max \varepsilon_j\}$. This event arises either when $\varepsilon_i = \bar{\varepsilon}$ or when all bidders are pessimistic so $\Pr A = p + (1-p)(1-p)^{n-1} = p(1+\rho)$. We have:

$$v_i(\gamma_i, \gamma) = (\sum_{n_0} \Pr(\tilde{n} = n_0, \varepsilon_i = \bar{\varepsilon}) \phi_{n_0}(\gamma_i - \gamma)(\gamma_i - \bar{\varepsilon})) + \Pr(A, \varepsilon_i = \underline{\varepsilon}) \phi_n(\gamma_i - \gamma)(\gamma_i - \underline{\varepsilon}).$$

Since $\phi'_n(0) = -1/\Delta$ for all $n > 1$, since $\phi'_1(0) = 0$, and since $\phi_{n_0}(0) = \frac{1}{n_0}$, the first order condition yields

$$\frac{1}{\Delta}[(\Pr A)(a - \bar{\varepsilon}) + (1-p)^n (a - \underline{\varepsilon})] \geq \sum \frac{1}{n_0} \Pr(A, \tilde{n} = n_0, \varepsilon_i = \bar{\varepsilon}) + \frac{1}{n} \Pr(A, \varepsilon_i = \underline{\varepsilon}).$$

The left hand side coincides with $\frac{\Pr A}{\Delta}[a - E[\varepsilon_i \mid A]$. Note also that the right hand side corresponds to the expected probability of winning for bidder i, so by symmetry it must equal $\frac{1}{n}$, and it can also be written as $(\Pr A) E[\frac{1}{\tilde{n}} \mid A]$, or equivalently (by symmetry again), $(\Pr A) E[\frac{1}{\tilde{n}}]$. So we obtain:

$$\gamma^*_\varepsilon \geq E[\varepsilon_i \mid A] + \Delta E[\frac{1}{\tilde{n}}]$$

or equivalently, since $\frac{1}{n} = (\Pr A) E[\frac{1}{\tilde{n}}]$, and since $\Pr A = p(1+\rho)$,

$$\gamma^*_\varepsilon \geq E[\varepsilon_i \mid A] + \frac{\Delta}{pn(1+\rho)}. \text{ QED}$$

[29] It is actually easy to check that $\mid \phi'_\varepsilon(0) \mid < 1/\Delta_\varepsilon$. Indeed, define $q(y) = \Pr(\varepsilon_1 - \varepsilon_2 > y)$ and denote by h the distribution over $x_2 - x_1$. We have:

$$\phi_\varepsilon(y) = \int q(y+z) h(z) dz.$$

Since the ε_i's are uniform, $\mid q'(y) \mid$ is maximum at $y = 0$, so $\mid \phi'_\varepsilon(0) \mid \leq \mid q'(0) \mid = 1/\Delta_\varepsilon$.

Proof of Proposition 13: Using $H(y)$ as defined in Proposition 10, and $h(y) = -\phi'_\varepsilon(y)$, we have:

$$v(\gamma_i, \gamma) = \int_{y \geq \gamma_i - \gamma} y h_\varepsilon(y) dx + \gamma \phi_\varepsilon(\gamma_i - \gamma) - H(\gamma_i - \gamma)$$

First order conditions thus yield $\gamma \phi'_\varepsilon(0) = H'(0)$, hence the desired result. QED

Proof of proposition 15: When the seller makes an offer equal to $p = v_1 + a_1$, the buyer accepts iff $\theta_2 - \theta_1 \geq a$, hence the seller obtains an expected payoff equal to

$$G_S = v_1 + a_1 \phi(a_1).$$

When the buyer makes an offer $p = v_2 - a_2$, the seller accepts if $p \geq v_1$, that is, if $\theta_2 - a_2 \geq \theta_1$, hence the buyer obtains an expected payoff equal to

$$G_B = a_2 \phi(a_2).$$

The optimal values of a_1 and a_2 are thus the same, and we call this value $a^* = \arg\max a\phi(a)$, and denote by G^*_S and G^*_B the corresponding gains for the seller and the buyer. Note that the expected surplus to be shared is the same whether the seller or the buyer makes the offer, and it is equal to $S(a^*)$. Who makes the offer thus only affects how the expected surplus is shared.

To see how the expected surplus $S(a^*)$ is shared, observe that when the buyer makes the offer, the seller obtains

$$\begin{aligned} R_S &= v_1 + E[\max(\theta_2 - \theta_1 - a^*, 0)] \\ &= v_1 + S(a^*) - a^* \phi(a^*). \end{aligned}$$

So the seller prefers to make the offer when $G^*_S > R_S$, that is, when

$$S(a^*) < 2a^* \phi(a^*).$$

Since $S(y) = \int_{x>y} -x\phi'(x)dx = y\phi(y) + \int_{x>y} \phi(x)dz$, we conclude the proof.

Under the split the difference mechanism the seller chooses a_1 and offers a price $p_1 = v_1 + a_1$, while the buyer chooses a_2 and offers a price $p_2 = v_2 - a_2$. The transaction takes place in the event $p_2 - p_1 \geq 0$, that is in the event $x_2 - x_1 \geq a_1 + a_2$, so the expected gain of the seller writes as:

$$\int_{y \geq a_1 + a_2} \frac{y + a_1 - a_2}{2} \phi(y) dy.$$

and similarly, the expected gain for the buyer can be written

$$\int_{y \geq a_1 + a_2} \frac{y + a_2 - a_1}{2} \phi(y) dy.$$

Let a^* as defined earlier. We verify that $a_1^* = a_2^* = a^*/2$ is an equilibrium. Assume 1 chooses $b_1 = a^*/2 + \delta$. Then he obtains a payoff $H(\delta)$ equal to:

$$H(\delta) = \frac{1}{2}\int_{y \geq a^*+\delta} (y+\delta)\phi(y)dy = \frac{1}{2}(\int_{y \geq a^*+\delta} (y-a^*)\phi(y)dy + G(a^*+\delta)).$$

Each of the terms on the right hand side is maximum for $\delta = 0$. So $a_1^* = a_2^* = a^*/2$ is an equilibrium.**QED**